

**A CONTEXT SWITCHING SYSTEM FOR A MULTI-THREAD EXECUTION
PIPELINE LOOP AND METHOD OF OPERATION THEREOF**

Inventors: Victor A. Bennett
5565 FM 549
Rockwall, Texas 75087

Sean W. McGee
9308 Bradner Drive
Austin, Texas 78748

Assignee: Agere Systems Guardian Corp.
9333 S. John Young Parkway
Room 301E1211
Orlando, Florida 32819

CERTIFICATE OF EXPRESS MAIL

I hereby certify that this correspondence, including the attachments listed, is being deposited with the United States Postal Service, Express Mail - Post Office to Addressee, Receipt No. EL 923683378 US, in an envelope addressed to Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the date shown below

2/25/02
Date of Mailing

EDITH SHEK
Typed or printed name of person mailing

Edith Shek
Signature of person mailing

Hitt Gaines & Boisbrun, P.C.
P.O. Box 832570
Richardson, Texas 75083
(972) 480-8800

A CONTEXT SWITCHING SYSTEM FOR A MULTI-THREAD EXECUTION PIPELINE LOOP AND METHOD OF OPERATION THEREOF

TECHNICAL FIELD OF THE INVENTION

[0001] The present invention is directed, in general, to network packet processors and, more specifically, to a context switching system for a multi-thread execution pipeline loop and method of operating the same.

BACKGROUND OF THE INVENTION

[0002] Communications networks are currently undergoing a revolution brought about by the increasing demand for real-time information being delivered to a diversity of locations employing multiple protocols. Many situations require the ability to transfer large amounts of data across geographical boundaries with increasing speed and accuracy. However, with the increasing size and complexity of the data that is currently being transferred, maintaining the speed and accuracy is becoming increasingly difficult.

[0003] Early communications networks resembled a hierarchical star topology. All access from remote sites was channeled back to a central location where a mainframe computer resided. Thus, each

transfer of data from one remote site to another, or from one remote site to the central location, had to be processed by the central location. This architecture is very processor-intensive and incurs higher bandwidth utilization for each transfer. This was not a major problem in the mid to late 1980s where fewer remote sites were coupled to the central location. Additionally, many of the remote sites were located in close proximity to the central location. Currently, hundreds of thousands of remote sites are positioned in various locations across assorted continents. Legacy networks of the past are currently unable to provide the data transfer speed and accuracy demanded in the marketplace of today.

[0004] In response to this exploding demand, data transfer through networks employing distributed processing has allowed larger packets of information to be accurately and quickly distributed across multiple geographic boundaries. Today, many communication sites have the intelligence and capability to communicate with many other sites, regardless of their location. This is typically accomplished on a peer level, rather than through a centralized topology, although a host computer at the central site can be appraised of what transactions take place and can maintain a database from which management reports are generated and operation issues addressed.

[0005] Distributed processing currently allows the centralized site to be relieved of many of the processor-intensive data

transfer requirements of the past. This is typically accomplished using a data network, which includes a collection of routers and/or switches. The routers and switches allow intelligent passing of information and data files between remote sites. However, increased demand and the sophistication required to route current information and data files, which may employ different protocols, quickly challenged the capabilities of existing routers and switches.

[0006] More specifically, network processors, such as the micro-processors employed in routers and switches, must be able to process multiple protocol data units (PDUs) at the same time. Typically, current network processors achieve multiprocessing of PDUs by assign an execution thread to each PDU. Each thread executes code, independently of the other threads, to process the PDUs. However, network processors are limited to specific amounts of memory on the chip or in cache memory to hold instructions and/or data. When a thread executing in the network processor needs to access off-chip memory, that thread is delayed until the request is fulfilled. The delay may cause the execution of all the other threads to be suspended until the request is fulfilled or prevent a new thread from being able to start execution. Another problem associated with the delay is the amount of precious thread execution cycles expended to process and determine if the request has been fulfilled. In view of the ever increasing demand for

higher transmission speeds these problems are highly undesirable.

[0007] Accordingly, what is needed in the art is a system to overcome the deficiencies of the prior art.

SUMMARY OF THE INVENTION

[0008] To address the above-discussed deficiencies of the prior art, the present invention provides a context switching system for a multi-thread execution pipeline loop having a pipeline latency and a method of operating the same. In one embodiment, the context switching system includes a context switch requesting subsystem configured to detect a device request from a thread executing within the multi-thread execution pipeline loop for access to a device having a fulfillment latency exceeding the pipeline latency, and generate a context switch request for the thread. Additionally, the context switching system includes a context controller subsystem configured to receive the context switch request and prevent the thread from executing until the device request is fulfilled.

[0009] In another embodiment, the present invention provides a method of operating a context switching system for use with a multi-thread execution pipeline loop having a pipeline latency, the method includes: (1) detecting a device request from a thread executing within the multi-thread execution pipeline loop for access to a device having a fulfillment latency exceeding the pipeline latency, (2) generating a context switch request for the thread when the thread issues the device request, and (3) receiving the context switch request and preventing the thread from executing

until the device request is fulfilled.

[0010] The present invention also provides, in one embodiment, a fast pattern processor that receives and processes protocol data units (PDUs) that includes a dynamic random access memory (DRAM) that contains instructions, a memory cache that caches certain of the instructions from the DRAM, and a tree engine that parses data within the PDUs and employs the DRAM and the memory cache to obtain ones of the instructions. The tree engine includes a multi-thread execution pipeline loop having a pipeline latency, and a context switching system for the multi-thread execution pipeline loop. The context switching system includes a context switch requesting subsystem that: (1) detects a device request from a thread executing within the multi-thread execution pipeline loop for access to a device having a fulfillment latency exceeding the pipeline latency, and (2) generates a context switch request for the thread when the thread issues the device request. The context switching system further includes a context controller subsystem that receives the context switch request and prevents the thread from executing until the device request is fulfilled.

[0011] The foregoing has outlined preferred and alternative features of the present invention so that those skilled in the art may better understand the detailed description of the invention that follows. Additional features of the invention will be described hereinafter that form the subject of the claims of the

invention. Those skilled in the art should appreciate that they can readily use the disclosed conception and specific embodiment as a basis for designing or modifying other structures for carrying out the same purposes of the present invention. Those skilled in the art should also realize that such equivalent constructions do not depart from the spirit and scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] For a more complete understanding of the present invention, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, in which:

[0013] FIGURE 1 illustrates a block diagram of an embodiment of a communications network constructed in accordance with the principles of the present invention;

[0014] FIGURE 2 illustrates a block diagram of an embodiment of a router architecture constructed in accordance with the principles of the present invention;

[0015] FIGURE 3 illustrates a block diagram of an embodiment of a fast pattern processor constructed in accordance with the principles of the present invention;

[0016] FIGURE 4 illustrates a block diagram of an embodiment of a pattern processing engine, generally designated 400, of a fast pattern processor constructed according to the principles of the present invention;

[0017] FIGURE 5 illustrates a block diagram of a context switching system for a multi-thread execution pipeline loop constructed according to the principles of the present invention; and

[0018] FIGURE 6 illustrates a flow diagram of an embodiment of a method of operating a context switching system for a multi-thread

execution pipeline loop constructed in accordance with the principles of the present invention.

DETAILED DESCRIPTION

[0019] Referring initially to FIGURE 1, illustrated is a block diagram of an embodiment of a communications network, generally designated 100, constructed in accordance with the principles of the present invention. The communications network 100 is generally designed to transmit information in the form of a data packet from one point in the network to another point in the network.

[0020] As illustrated, the communications network 100 includes a packet network 110, a public switched telephone network (PSTN) 115, a source device 120 and a destination device 130. In the illustrative embodiment shown in FIGURE 1, the packet network 110 comprises an Asynchronous Transfer Mode (ATM) network. However, one skilled in the art readily understands that the present invention may use any type of packet network. The packet network 110 includes routers 140, 145, 150, 160, 165, 170 and a gateway 155. One skilled in the pertinent art understands that the packet network 110 may include any number of routers and gateways.

[0021] The source device 120 may generate a data packet to be sent to the destination device 130 through the packet network 110. In the illustrated example, the source device 120 initially sends the data packet to the first router 140. The first router 140 then determines from the data packet which router to send the data packet to based upon routing information and network loading. Some

information in determining the selection of a next router may include the size of the data packet, loading of the communications link to a router and the destination. In this example, the first router 140 may send the data packet to the second router 145 or fourth router 160.

[0022] The data packet traverses from router to router within the packet network 110 until it reaches the gateway 155. In one particular example, the data packet may traverse along a path that includes the first router 140, the fourth router 160, the fifth router 165, the sixth router 170, the third router 150 and finally to the gateway 155. The gateway 155 converts the data packet from the protocol associated with the packet network 110 to a different protocol compatible with the PSTN 115. The gateway 155 then transmits the data packet to the destination device 130 via the PSTN 115. However, in another example, the data packet may traverse along a different path such as the first router 140, the second router 145, the third router 150 and finally to the gateway 155. It is generally desired when choosing a subsequent router, the path the data packet traverses should result in the fastest throughput for the data packet. It should be noted, however, that this path does not always include the least number of routers.

[0023] Turning now to FIGURE 2, illustrated is a block diagram of an embodiment of a router architecture, generally designated 200, constructed in accordance with the principles of the present

invention. The router architecture 200, in one embodiment, may be employed in any of the routers illustrated in FIGURE 1. The router architecture 200 provides a unique hardware and software combination that delivers high-speed processing for multiple communication protocols with full programmability. The unique combination provides the programmability of traditional reduced instruction set computing (RISC) processors with the speed that, until now, only application-specific integrated circuit (ASIC) processors could deliver.

[0024] In the embodiment shown in FIGURE 2, the router architecture 200 includes a physical interface 210, a fast pattern processor (FPP) 220, a routing switch processor (RSP) 230, and a system interface processor (SIP) 240. The router architecture 200 may also include a fabric interface controller 250 which is coupled to the RSP 230 and a fabric network 260. It should be noted that other components not shown may be included within the router architecture 200 without departing from the scope of the present invention.

[0025] The physical interface 210 provides coupling to an external network. In an exemplary embodiment, the physical interface 210 is a POS-PHY/UTOPIA level 3 interface. The FPP 220, in one embodiment, may be coupled to the physical interface 210 and receives a data stream that includes protocol data units from the physical interface 210. The FPP 220 analyzes and classifies the

protocol data units and subsequently concludes processing by outputting packets to the RSP 230.

[0026] The FPP 220, in conjunction with a powerful high-level functional programming language (FPL), is capable of implementing complex pattern or signature recognition and operates on the processing blocks containing those signatures. The FPP 220 has the ability to perform pattern analysis on every byte of the payload plus headers of a data stream. The pattern analysis conclusions may then be made available to a system logic or to the RSP 230, allowing processing block manipulation and queuing functions. The FPP 220 and RSP 230 provide a solution for switching and routing. The FPP 220 further provides glueless interfaces to the RSP 230 and the SIP 240 to provide a complete solution for wire-speed processing in next-generation, terabit switches and routers.

[0027] As illustrated in FIGURE 2, the FPP 220 employs a first communication link 270 to receive the data stream from the physical interface 210. The first communication link 270 may be an industry-standard UTOPIA Level 3/UTOPIA Level 2/POS-PHY Level 3 interface. Additionally, the FPP 220 employs a second communication link 272 to transmit packet and conclusions to the RSP 230. The second communication link 272 may be a POS-PHY Level 3 interface.

[0028] The FPP 220 also includes a management path interface (MPI) 275, a function bus interface (FBI) 280 and a configuration

bus interface (CBI) 285. The MPI 275 enables the FPP 220 to receive management frames from a local microprocessor. In an exemplary embodiment, this may be handled through the SIP 240. The FBI 280 connects the FPP 220 and the SIP 240, or custom logic in certain situations, for external processing of function calls. The CBI 285 connects the FPP 220 and other devices (e.g., physical interface 210 and RSP 230) to the SIP 240. Other interfaces (not shown), such as memory interfaces, are also well within the scope of the present invention.

[0029] The FPP 220 provides an additional benefit in that it is programmable to provide flexibility in optimizing performance for a wide variety of applications and protocols. Because the FPP is a programmable processor rather than a fixed-function ASIC, it can handle new protocols or applications as they are developed as well as new network functions as required. The FPP 220 may also accommodate a variety of search algorithms. These search algorithms may be applied to large lists beneficially.

[0030] The RSP 230 is also programmable and works in concert with the FPP 220 to process the protocol data units classified by the FPP 220. The RSP 230 uses the classification information received from the FPP 220 to determine the starting offset and the length of the Protocol data unit payload, which provides the classification conclusion for the Protocol data unit. The classification information may be used to determine the port and

the associated RSP 230 selected for the Protocol data unit. The RSP 230 may also receive additional Protocol data unit information passed in the form of flags for further processing.

[0031] The RSP 230 also provides programmable traffic management including policies such as random early discard (RED), weighted random early discard (WRED), early packet discard (EPD) and partial packet discard (PPD). The RSP 230 may also provide programmable traffic shaping, including programmable per queue quality of service (QoS) and class of service (CoS) parameters. The QoS parameters include constant bit rate (CBR), unspecified bit rate (UBR), and variable bitrate (VBR). Correspondingly, CoS parameters include fixed priority, round robin, weighted round robin (WRR), weighted fair queuing (WFQ) and guaranteed frame rate (GFR).

[0032] Alternatively, the RSP 230 may provide programmable packet modifications, including adding or stripping headers and trailers, rewriting or modifying contents, adding tags and updating checksums and CRCs. The RSP 230 may be programmed using a scripting language with semantics similar to the C language. Such script languages are well known in the art. Also connected to the RSP 230 are the fabric interface controller 250 and the fabric network 260. The fabric interface controller 250 provide the physical interface to the fabric 260, which is typically a communications network.

[0033] The SIP 240 allows centralized initialization and

configuration of the FPP 220, the RSP 230 and the physical interfaces 210, 250. The SIP 240, in one embodiment, may provide policing, manage state information and provide a peripheral component interconnect (PCI) connection to a host computer. The SIP 240 may be a PayloadPlus™ Agere System Interface commercially available from Agere Systems, Inc.

[0034] Turning now to FIGURE 3, illustrated is a block diagram of an embodiment of a fast pattern processor (FPP), generally designated 300, constructed in accordance with the principles of the present invention. The FPP 300 includes an input framer 302 that receives protocol data units via external input data streams 330, 332. The input framer 302 frames packets containing the protocol data units into 64-byte processing blocks and stores the processing blocks into an external data buffer 340. The input data streams 330, 332 may be 32-bit UTOPIA/POS-PHY from PHY and 8-bit POS-PHY management path interface from SIP 240 (FIGURE 2), respectively.

[0035] Typically, a data buffer controller 304 is employed to store the processing blocks to the external data buffer 340. The data buffer controller 304 also stores the processing blocks and associated configuration information into a portion of a context memory subsystem 308 associated with a context, which is a processing thread. As illustrated, the context memory subsystem 308 is coupled to a data buffer controller 304.

[0036] Additionally, the context memory subsystem 308 is coupled to a checksum/cyclical redundancy check (CRC) engine 314 and a pattern processing engine 312. The checksum/CRC engine 314 performs checksum or CRC functions on processing block and on the protocol data units embodied with the processing block. The pattern processing engine 312 performs pattern matching to determine how protocol data units are classified and processed. The pattern processing engine 312 is coupled to a program memory 350.

[0037] The FPP 300 further includes a queue engine 316 and an arithmetic logic unit (ALU) 318. The queue engine 316 manages replay contexts for the FPP 300, provides addresses for block buffers and maintains information on blocks, protocol data units, and connection queues. The queue engine 316 is coupled to an external control memory 360 and the internal function bus 310. The ALU 318 is coupled to the internal function bus 310 and is capable of performing associated computational functions.

[0038] Also coupled to the internal function bus 310 is a functional bus interface 322. The functional bus interface 322 passes external functional programming language function calls to external logic through a data port 336. In one exemplary embodiment, the data port 336 is a 32-bit connection to the SIP 240 (FIGURE 2). The FPP 300 also includes a configuration bus interface 320 for processing configuration requests from externally

coupled processors. As illustrated, the configuration bus interface 320 may be coupled to a data port 334, such as an 8-bit CBI source.

[0039] Additionally, coupled to the internal function bus 310 is an output interface 306. The output interface 306 sends protocol data units and their classification conclusions to the downstream logic. The output interface 306 may retrieve the processing blocks stored in the data buffer 340 and send the protocol data units embodied within the processing blocks to an external unit through an output data port 338. The output data port 338, in an exemplary embodiment, is a 32-bit POS-PHY connected to the RSP 230 (FIGURE 2). Additional background information concerning the FPP is discussed in U.S. Patent Application Serial No. 9/798,472, titled "A VIRTUAL REASSEMBLY SYSTEM AND METHOD OF OPERATION THEREOF," and is incorporated herein by reference as if reproduced herein in its entirety.

[0040] Turning now to FIGURE 4, illustrated is a block diagram of an embodiment of a pattern processing engine, generally designated 400, of a fast pattern processor constructed according to the principles of the present invention. The pattern processing engine, in one embodiment, is similar to the pattern processing engine 312 of FIGURE 3 and performs pattern matching to determine how the protocol data units (PDUs) are classified and processed.

[0041] In the illustrated embodiment, the pattern processing

engine 400 includes first and second flow engines 402, 404, a first-in-first-out buffer (FIFO) 410 and a tree engine 420. The pattern processing engine 400 is also coupled to a memory cache 430 and a dynamic random access memory (DRAM) 440. The tree engine 420 includes a multi-thread execution pipeline loop 422 and a context switching system 424. The tree engine is also coupled to the memory cache 430 and indirectly to the DRAM 440.

[0042] The pattern processing engine 400 employs the first and second flow engines 402, 404 to process the processing blocks based on their associated contexts. As described previously, the PDUs are framed into processing blocks for processing. Also, the processing of a PDU has an associated context, which is a processing thread (thread). Each of the first and second flow engines 402, 404 operate in a parallel, pipeline manner and are configured to process at least one of the processing blocks based on its context. In the illustrated embodiment, the first flow engine 402 processes even number contexts and the second flow engine 404 processes odd number contexts. Typically, the first and second flow engines 402, 404 may have several processing blocks associated with several contexts that are being processed at any one time. When the first and second flow engines 402, 404 finish processing a context, the first and second flow engines 402, 404 place the finished context (or thread) in the FIFO 410 to await processing by the tree engine 420.

[0043] In a related embodiment, the first and second flow engines 402, 404 transfer context (or thread) to the tree engine 420 when the FPL code associated with that context makes a call to be processed by the tree engine 420 to perform one or more functions, such as pattern matching. When the tree engine 420 finishes performing the function or functions, the tree engine 420 returns a result to the appropriate flow engine. That flow engine then resumes processing that context.

[0044] The tree engine 420, in one embodiment, is configured to process multiple contexts (of threads) that employ function trees to perform pattern matching or data validation on the PDU data contained within specific processing blocks or on at least a portion of the processing blocks. Function trees are a set of functions arranged in a tree structure. Each function tree has a root and can have any number branches off of the root. Each branch may also have any number of sub-branches and so on. Function tree processing starts at a root function and the outcome of the root function determines which branch to take. Each branch performs another function and the outcome of which determines the next branch to take and so on. One skilled in the art is familiar with tree structures having multiple branches. Also, for purposes of the present invention, the phrase "configured to" means that the device, the system or the subsystem includes the necessary software, hardware, firmware or a combination thereof to accomplish

the stated task.

[0045] The tree engine 420 also employs the multi-thread execution pipeline loop 422 to sequence each thread through execution of its associated function tree. The multi-thread execution pipeline loop 422 has a number of stages, where each stage maintains information for a thread currently executing. See FIGURE 5 for more information concerning a multi-thread execution pipeline loop and its associated stages. The tree engine 420 retrieves a thread from the FIFO 410 and places that thread at the beginning of the multi-thread execution pipeline loop 422 to start executing. As each thread traverses the multi-thread execution pipeline loop 422, the threads perform at least one function of its associate function tree on at least a portion of one or more of the associated processing blocks. In this manner, each thread sequences through the data depending upon the outcome of the function performed.

[0046] For example, one thread may first match the first three bits of the data. If the match was successful, then at the next stage of the multi-thread execution pipeline loop 422, the thread may try to match only the next two bits of the data. The tree engine 420 also allows each thread to take different branches depending upon the outcome of the function performed. If a thread reaches the end of the multi-thread execution pipeline loop 422 and the thread has not completed its processing, the thread may be

looped back to the beginning of the multi-thread execution pipeline loop 422 to continue its processing. If a thread has finished its processing, the thread may return to one of the first or second flow engines 402, 404 or to another portion of the fast pattern processor for additional processing. In another embodiment, the thread may exit the multi-thread execution pipeline loop 422 at any stage without having to sequence to the end of the multi-thread execution pipeline loop 422.

[0047] Since the tree engine 420 may have a number of threads in its multi-thread execution pipeline loop 422 and each function tree can contain any number of functions and branches, the tree engine 420 employs the memory cache 430 and the DRAM 440 to retrieve the current function of the function tree for each thread. If the current function for a particular thread is not in the memory cache 430, a request to retrieve the current function is made to the DRAM 440. Fulfilling requests from the DRAM 440 typically incurs a longer fulfillment time than does fulfilling requests from the memory cache 430. This longer fulfillment time may be longer than the pipeline latency and cause the delay of one or more of the threads or prevent a new thread from being able to be added to the multi-thread execution pipeline loop 422 to begin execution. For purposes of the present invention, "pipeline latency" is the rate at which data or a thread traverses a multi-thread execution pipeline loop.

[0048] In the illustrated embodiment, the present invention advantageously employs the context switching system 424 to manage fulfillment requests that exceed the pipeline latency of the multi-thread execution pipeline loop 422. The context switching system 424 is configured to detect a device request from a thread executing within the multi-thread execution pipeline loop 422 for access to a device having a fulfillment latency that exceeds the pipeline latency and switches context for that thread, thus preventing that thread from executing until the device request is fulfilled. For purposes of the present invention, "fulfillment latency" is the rate at which it takes a device to fulfill a request. Also, see FIGURE 5 for a more detailed description of the context switching system.

[0049] For example, if a thread within the multi-thread execution pipeline loop 422 issues a device request to access data in the memory cache 430 and the data is not within the memory cache 430, then a device request is made to obtain the data from the DRAM 440. Since the time to fulfill the device request from the DRAM 440 is longer than the pipeline latency, the context switching system 424 will prevent that particular thread from executing until the device request from the DRAM 440 is fulfilled. Thus, the context switching system 424 allows the other threads to continue to execute and does not cause the waste of execution cycles for that particular thread. The context switching system 424, in one

embodiment, allows the thread that was prevented from executing to continue to traverse the multi-thread execution pipeline loop 422 until the device request is fulfilled. In this example, the thread would continue to traverse the multi-thread execution pipeline loop 422 until the DRAM 440 fulfilled the request for the data requested. Then, that thread would continue processing with another function from its associated function tree.

[0050] Turning now to FIGURE 5, illustrated is a block diagram of a context switching system, generally designated 520, for a multi-thread execution pipeline loop 500 constructed according to the principles of the present invention. The multi-thread execution pipeline loop 500 may be used to sequence a thread (or a context) through its execution. In one embodiment, the multi-thread execution pipeline loop 500 may be used to sequence threads through their associated function trees. See FIGURE 4 for a discussion of function trees.

[0051] In the illustrated embodiment, the multi-thread execution pipeline loop 500 includes 10 pipeline stages 502. Each of the pipeline stages 502 maintains information for the thread currently executing in that particular pipeline stage. Of course, however, the multi-thread execution pipeline loop 500 is not limited to only 10 pipeline stages and may have any number of pipeline stages depending upon its particular implementation.

[0052] The multi-thread execution pipeline loop 500 receives a

thread to process through a receive line 510 and stores the thread's information at the beginning pipeline stage 504 when the beginning pipeline stage 504 is empty. In another embodiment, a new thread may be stored in any of the pipeline stages 502 that are empty. As the multi-thread execution pipeline loop 500 sequences, each thread moves to the next pipeline stage 502 and performs another function. When a thread reaches the end or last pipeline stage 506 of the multi-thread execution pipeline loop 500 and the thread has not finished processing, that thread is looped back to the beginning pipeline stage 504 of the multi-thread execution pipeline loop 500. If the thread has finished processing, the thread is sent out the output line 540. The finished thread may be sent to another processor, sub-processor or another area for additional processing. In a related embodiment, when a thread finishes processing, the thread may be removed from its current pipeline stage 502 and not wait until it reaches the last pipeline stage 506 of the multi-thread execution pipeline loop 500 to be removed.

[0053] The multi-thread execution pipeline loop 500 also has an associated pipeline latency. As described previously, a pipeline latency is the rate at which data or a thread traverses the multi-thread execution pipeline loop. For example, each pipeline stage may allow two clock cycles of execution for a thread in any given pipeline stage. Thus, multi-thread execution pipeline loop 500 has a pipeline latency of two clock cycles. In another embodiment,

another method of defining a pipeline latency is the number of clock cycles for all of the pipeline stages of the multi-thread execution pipeline loop 500. For example, the illustrated multi-thread execution pipeline loop 500 includes 10 pipeline stages 502. If each pipeline stage allows two clock cycles of execution time, the pipeline latency for all of the pipeline stages would be 20 clock cycles. Of course, however, other methods of defining a pipeline latency are well within the scope of the present invention.

[0054] Associated with the multi-thread execution pipeline loop 500 is a memory device 530. The memory device 530 include a memory cache 532 coupled to a dynamic random access memory (DRAM) 534. The DRAM 534 may contain instructions for the threads executing within the multi-thread execution pipeline loop 500. The DRAM 534 may also contain data. The memory cache 532 caches certain instructions of the DRAM 534. In another embodiment of the present invention, the memory cache 532 may cache certain instructions, data or a combination thereof of the DRAM 534. The cache memory 532 may be conventional cache memory that is local or within the same processor as the multi-thread execution pipeline loop 500 and the DRAM 534 may be conventional external DRAM. However, the amount of the memory cache 532 available is typically smaller than the DRAM 534 due to the higher cost and limited space availability. Also, the memory cache 532 generally has a smaller fulfillment

latency than the fulfillment latency of the DRAM 534. The longer fulfillment latency of the DRAM 534 is typically due to accessing an external memory device.

[0055] The multi-thread execution pipeline loop 500 is also coupled to the context switching system 520. The context switching system 520, in one embodiment, includes a context switch requesting subsystem 522 and a context controller subsystem 524. The context switch requesting subsystem 522 is configured to detect a device request from a thread executing within the multi-thread execution pipeline loop 500 for access to a device having a fulfillment latency exceeding the pipeline latency of the multi-thread execution pipeline loop 500. The context switch requesting subsystem 522 is further configured to generate a context switch request for the thread that issued the device request.

[0056] The context controller subsystem 524 is configured to receive the context switch request from the context switch requesting subsystem 522 and prevent the thread from executing until the device request is fulfilled. The context controller subsystem 524, in one embodiment, is further configured to replace the thread's current instruction with a NO-Operation (NOP) instruction to prevent the thread from executing until the device request is fulfilled. The context controller subsystem 524 may also allow the thread to continue to traverse the multi-thread execution pipeline loop 500 while waiting for the device request to

be fulfilled. In a related embodiment, the context controller subsystem 524 is further configured to allow the other threads within the multi-thread execution pipeline loop 500 to continue to execute while the thread that made the device request is waiting for the device request to be fulfilled.

[0057] In one embodiment, the device request may be a request to access external memory due to a cache miss status. For example, a thread may request to access an instruction from the memory cache 532. If the information is not currently in the memory cache 532, a cache miss status is issued by the memory device 530 and a request is made to the DRAM 534 to access the desired information. As stated previously, the DRAM 534 has a longer fulfillment latency and is typically longer than the pipeline latency of the multi-thread execution pipeline loop 500. If the multi-thread execution pipeline loop 500 delayed execution of all the threads within the pipeline until the request for one thread to the DRAM 534 was fulfilled, the throughput of the multi-thread execution pipeline loop 500 would be unduly delayed. Also, the multi-thread execution pipeline loop 500 may not be able to maintain a desired processing bandwidth.

[0058] The present invention, in one embodiment, advantageously overcomes the problems associated with a device request to access a device having a fulfillment latency exceeding the pipeline latency of the multi-thread execution pipeline loop 500 by

employing the context switching system 520. In the above example, the context switch requesting subsystem 522 or the context switching system 520 detects the device request to access external memory, such as the DRAM 534, due to a cache miss status. Upon detecting the device request, the context switch requesting subsystem 522 also generates a context switch request for the thread that issued the device request. The context controller subsystem 524 of the context switching system 520 receives the context switch request and prevents that thread from executing until the device request to the external memory is fulfilled. The context controller subsystem 524, in one embodiment, also allows the other threads within the multi-thread execution pipeline loop 500 to continue to execute. Thus, the context switching system 520 can maintain the throughput of the multi-thread execution pipeline loop 500 and maintain the desired processing bandwidth.

[0059] The context switching system 520 may further include a miss fulfillment first-in-first-out buffer ("miss fulfillment FIFO") 550 to accommodate a thread or threads that are waiting for fulfillment of device requests to devices having a fulfillment latency that exceeds the pipeline latency of the multi-thread execution pipeline loop 500. In one embodiment, the context controller subsystem 524 is further configured to employ the miss fulfillment FIFO 550 to store the thread and/or its related information in the miss fulfillment FIFO 550 upon reaching the end

position (or the last pipeline stage 506) of the multi-thread execution pipeline loop 500. In a related embodiment, the context controller subsystem 524 may store the thread in the miss fulfillment FIFO 550 upon receiving a context switch request for that thread instead of waiting for the thread to reach the end position of the multi-thread execution pipeline loop 500. By the context controller subsystem 524 storing the thread in the miss fulfillment FIFO 550 upon receiving a context switch request, the context switching system 520 advantageously allows the multi-thread execution pipeline loop 500 to receive and process a new thread. Thus, the multi-thread execution pipeline loop 500 may be filled with threads performing useful work and, as such, increasing the processing throughput.

[0060] The context controller subsystem 524 is further configured to sequence the stored thread through the miss fulfillment FIFO 550 and reinsert the thread into the multi-thread execution pipeline loop 500 at a beginning position (or the beginning pipeline stage 504). In a related embodiment, the context controller subsystem 524 is further configured to sequence the thread through the miss fulfillment FIFO 550 at a rate equal to the pipeline latency of the multi-thread execution pipeline loop 500. Of course, however, the present invention is not limited to sequencing at a rate equal to the pipeline latency. Other embodiments of the present invention may sequence the thread

through the miss fulfillment FIFO 550 at any rate.

[0061] Once a thread has been stored in the miss fulfillment FIFO 550, a new thread may be inserted at the beginning pipeline stage 504 of the multi-thread execution pipeline loop 500 when the beginning pipeline stage 504 is an empty. Alternatively, a new thread may be stored in any empty pipeline stage. Thus, the context switching system 520 advantageously employs the miss fulfillment FIFO 550 to allow the thread to wait for its device request to be fulfilled or delay the thread for a period of time, while allowing another thread to start executing in the multi-thread execution pipeline loop 500.

[0062] One skilled in the art should know that the present invention is not limited to switching context based upon a request to DRAM that has a fulfillment latency greater than the pipeline latency. In other embodiments, the present invention may perform context switching for any request to any device that has a fulfillment latency that exceeds the pipeline latency.

[0063] Turning now to FIGURE 6, illustrated is a method of operating a context switching system, generally designated 600, for a multi-thread execution pipeline loop constructed according to the principles of the present invention. In FIGURE 6, the method 600 first performs initialization in a step 602.

[0064] After initialization, the method 600 determines if there is a device request from a thread executing within the multi-thread

execution pipeline loop for access to a device in a decisional step 604. If the method 600 determined that there was a device request, the method then determines if the device request is to a device having a fulfillment latency exceeding the pipeline latency of the multi-thread execution pipeline loop in a decisional step 610. If the device does not have a fulfillment latency that exceeds the pipeline latency, the method 600 allows the device request to process as normal and returns to determine the next type of request in the decisional step 604. If the device has a latency that does exceed the pipeline latency, the method 600 then generates a context switch request for that thread in a step 620. For example, the device request may be to access an external DRAM, which typically has a fulfillment latency that exceeds the pipeline latency. The method 600 then returns to determine the next type of request in the decisional step 604.

[0065] If in the method 600 did not have a device request in the decisional step 604, the method 600 then determines if it received a context switch request in a decisional step 630. If a context switch request was received, the method 600 then prevents the thread associated with the contact switch request from executing until the device request is fulfilled in a step 640. In a related embodiment, preventing the thread from executing may further include replacing the thread's current instruction with a NOP instruction to prevent the thread from executing until the device

request is fulfilled. Of course, however, the present invention is not limited to using NOP instructions to prevent a thread from executing. In other embodiments of the present invention, any type of instruction or flag may be used to prevent the thread from executing. In another embodiment, the method 600 may also allow the threads within the multi-thread execution pipeline loop to continue to execute while the thread is waiting for its device request to be fulfilled. Next, the method 600 returns to determine the next type of request in the decisional step 604.

[0066] If a context switch request was not received in the decisional step 630, the method 600 then determines if the device request for the associated thread has been fulfilled in a decisional step 650. If the device request has been fulfilled, the method 600 then sets the thread to allow execution again in a step 660. In one embodiment, the method 600 may replace the thread's NOP instruction with the thread's original instruction before the device request was issued. In a related embodiment, the method 600 may store a retrieved instruction from the device and then allow the thread to execute the retrieved instruction. The method 600 then returns to determine the next type of request in the decisional step 604.

[0067] If it was determined that the device request was not fulfilled in the decisional step 650, the method 600 then sequences the thread that was prevented from executing through the multi-

thread execution pipeline loop in a step 670. In another embodiment, the method 600 allows the other threads to also sequence through the multi-thread execution pipeline loop.

[0068] In the illustrated embodiment, the method 600 then determines if the thread that was prevented from executing has reached the end (or last pipeline stage) of the multi-thread execution pipeline loop in a decisional step 680. If the thread did not reach the end of the multi-thread execution pipeline loop, the method 600 then returns to determine the next type of request in the decisional step 604. If the thread has reached the end of the multi-thread execution pipeline loop, the method may store the thread in a miss fulfillment FIFO in a step 690. The method 600 may also sequence the thread through the miss fulfillment FIFO at a rate equal to the pipeline latency of the multi-thread execution pipeline loop. Of course, however, the present invention may sequence the thread through the miss fulfillment FIFO at any rate. Once the thread reaches the end of the miss fulfillment FIFO, the thread is reinserted at the beginning of the multi-thread execution pipeline loop to continue processing. The method 600 then returns to determine the next type of request in the decisional step 604.

[0069] One skilled in the art should know that the present invention is not limited to processing only one device request from a thread for access to a device having a fulfillment latency exceeding the pipeline latency of the multi-thread execution

pipeline loop. The present invention and method may process any number of device requests from any number of threads for access to devices having a fulfillment latency exceeding the pipeline latency of the multi-thread execution pipeline loop. Also, other embodiments of the present invention may have additional or fewer steps than described above.

[0070] Although the present invention has been described in detail, those skilled in the art should understand that they can make various changes, substitutions and alterations herein without departing from the spirit and scope of the invention in its broadest form.